Can You Trust Autonomous Vehicles: Contactless Attacks against Sensors of Self-Driving Vehicles

Chen Yan
Zhejiang University
yanchen@zju.edu.cn

Wenyuan Xu
Zhejiang University
& University of South Carolina
wyxu@cse.sc.edu

Jianhao Liu
Qihoo 360
liujianhao@360.cn

ABSTRACT

To improve road safety and driving experiences, autonomous vehicles have emerged recently, and they can sense their surroundings and navigate without human intervention. Although promising and improving safety features, the trustworthiness of these cars has to be examined before they can be widely adopted on the road. Unlike traditional network security, autonomous vehicles rely heavily on their sensory ability of their surroundings to make driving decision, which makes sensors an interface for attacks. Thus, in this paper we examine the security of the sensors of autonomous vehicles, and investigate the trustworthiness of the 'eyes' of the cars.

Our work investigates sensors whose measurements are used to guide driving, i.e., millimeter-wave radars, ultrasonic sensors, forward-looking cameras. In particular, we present contactless attacks on these sensors and show our results collected both in the lab and outdoors on a Tesla Model S automobile. We show that using off-the-shelf hardware, we are able to perform jamming and spoofing attacks, which caused the Tesla's blindness and malfunction, all of which could potentially lead to crashes and impair the safety of self-driving cars. To alleviate the issues, we propose software and hardware countermeasures that will improve sensor resilience against these attacks.

Keywords

Autonomous vehicles; security; ultrasonic sensors; millimeter-wave radars; cameras

1. INTRODUCTION

Improving road safety, driving experiences, and driving efficiency has long been a focus of the automotive industry, and already we have witnessed the rapid development of Advanced Driver Assistance Systems (ADAS), which can sense its driving environment and warn drivers of immediate dangers. With the advances in sensing technology and information fusion, vehicles are going forward into a new era — fully autonomous vehicles. Numerous major companies and research organizations have developed their prototype autonomous cars. For instance, Tesla Motors has popularized driverless technology with its Autopilot system.

The safety of autonomous cars has been a focus of the prolonged debate over this technology. Comparing to traditional ones, autonomous vehicles requires almost no human inputs for driving control, therefore safety relies purely on the on-board computing systems, which in turn depend on sensors and their measurements of the surroundings to make driving decisions. Being the 'eyes' of on-board computing systems, sensors play an important role in autonomous vehicle safety, and their accuracy and immediacy have to be guaranteed to achieve safe autonomous driving.

The industry has been working on improving the accuracy and robustness of sensors. Nevertheless, recently a Tesla Model S car crashed into a white truck and caused one death while the driver fully relied on the Autopilot system [25]. The accident indicates that existing sensors cannot reliably detect neighboring cars even in normal yet special road conditions, not to mention intentional attacks against these sensors. In light of the fact that the security issues of sensors have not earned their due attention, we investigate attacks that utilize the underlying principles of sensors to blind or deceive them, e.g., exploiting the active probing mechanisms that are used to detect barriers. This type of attacks can lead to malfunctions, falsified readings, or even physical damage, and the consequences could be fatal both to one car and to a collection of cars nearby, i.e., in a Vehicle to Vehicle (V2V) network.

Understanding the attack methods, its feasibility, and its influences on sensor readings as well as autonomous car behaviors will provide insights for improving the safety of self-driving automobiles. In this work, we performed an empirical security study on the sensors of autonomous cars. Specifically, we studied and examined three types of essential automotive sensors that are widely used for autonomous driving, i.e., ultrasonic sensors, Millimeter Wave Radars, and cameras. We have carried out several attacks against them, and proved the destructive impact of attacks on the sensor data, as well as on the automated driving systems by experiments on a Tesla Model S sedan.

We summarize our contributions as follows.

- We raise the security risks and concerns of sensors used for Automated Driving and Advanced Driver Assistance Systems.
- To the best of our knowledge, we are the first to experimentally examine the feasibility of launching contactless attacks on automotive ultrasonic sensors and MMW Radars. Our experiments in the laboratory and outdoors on vehicles have demonstrated the consequences of jamming and spoofing attacks by exploiting the underlying sensing principles.

- We have verified the attacks on four different cars including a Tesla Model S with the Autopilot system, and demonstrated the impact of these attacks on automated driving system.

The rest of this paper is organized as follows. Background and related work on vehicle security are given in Section 2. We introduce automated driving system and relevant sensors in Section 3, and discuss the threat model and study procedure in Section 4. The details of attacks on ultrasonic sensors, MMW Radars, and cameras are given in Section 5, 6, and 7, respectively. In Section **??** we discuss the attack feasibility and countermeasures, as well as limitations and future work. Section 8 concludes the paper.

2. BACKGROUND AND RELATED WORK

The security of automotive systems has been studied for more than a decade. It is known that the security risk stems from the structure of automotive system, i.e., the interconnection of communication buses and *Electronic Control Units* (ECUs). Today, the infrastructure of modern vehicles is designed in such a way that all components are networked with each other by the CAN-bus, and they can exchange data as well as control commands via the bus. This structure facilitates the functionality and efficiency of modern vehicles, but poses a serious threat in addition to potential insecure components [31, 32]. For example, security breach on one ECU (especially those with external connections, e.g., telematics) could possibly lead to the exploitation of other safety-critical ECUs through the unprotected bus network (e.g., CAN bus) and endangers the whole vehicle.

Several studies [12, 27] have shown the feasibility of launching CAN-bus attacks, mainly through OBD-II port, to cause malfunction and even take control of the car. It has been demonstrated that an attacker who is able to infiltrate virtually any ECUs can leverage this ability to completely circumvent a broad array of safety-critical systems, such as falsifying the control panel displays, disabling the brakes, killing the engine, and rolling the steering wheel.

In addition, it is possible to launch attacks without any physical access to a car. Checkoway et al. [3] analyzed the external attack surfaces of a modern automobile, and discovered that remote exploitation is feasible via a broad range of attack vectors (including mechanics tools, CD players, Bluetooth and cellular radio), and further, that wireless communications channels allow long distance vehicle control, location tracking, in-cabin audio exfiltration and theft. Miller and Valasek, after their survey [15] of 21 popular car models, performed a remote attack against un unaltered Jeep Cherokee that resulted in physical control of the vehicle [16].

Previous researches on vehicle security mostly focused on the internal network and Electronic Control Units (e.g., telematics and immobilizer). However, few attention has been devoted to sensors. Existing attacks depend mainly on vulnerable information interfaces, while the sensory (physical) channels have not attracted their due attention and shall be exploited thoroughly.

Petit et al. has recently raised people's attention to sensors by studying LiDAR and cameras [19]. Their work focused on remote attacks on camera-based system and LiDAR using commodity hardware, which achieved effective blinding, jamming, replay, relay, and spoofing attacks.

In our research, we focus on the security of popular vehicular sensors that have already been widely used in *Advanced Driver Assistance System* (ADAS) and self-driving cars. We will show experiment results that were conducted both in laboratories and on popular cars, including models of Tesla, Audi, Volkswagen, and Ford.

3. SYSTEM OVERVIEW

In this section we give a brief introduction to the sensor technologies that Automated Driving Systems and Advanced Driver Assistance Systems are based on, and discuss the motivation to examine ultrasonic sensors, Millimeter Wave (MMW) Radars, and cameras.

3.1 Sensor Overview

Before discussing the detailed principles underlying these sensors, we overview their features and compare their differences.

Sensor categories. Ultrasonic sensors, MMW radars, cameras, and LiDAR are key sensors on current self-driving vehicles. Each is designed for its dedicated sensing range. Nevertheless, they, in combination, can detect obstacles in a wide range. They can be roughly divided into proximity, short-range, medium-range, and long-range, as shown in Figure 1. Note that MMW radars can be designed to cover a sensing range of short-range, medium-range, or long-range, by selecting the appropriate frequency bands (e.g., 24 GHz or 77 GHz), antenna design, and etc.

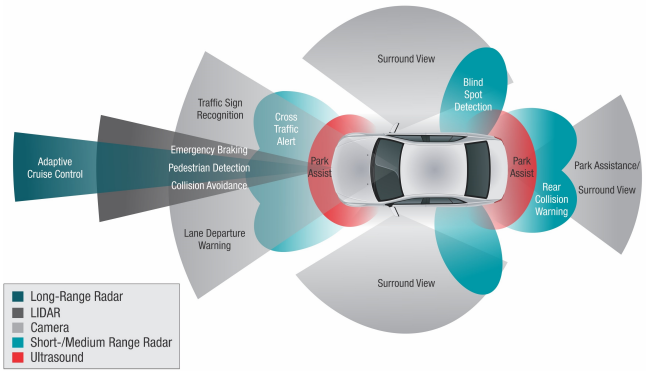

Figure 1: Major ADAS sensor types and typical vehicle positions [23].

1. *Proximity* (2 m). Ultrasonic sensors are proximity sensors that aim at detecting barriers within several meters from a car body. They are mainly designed for low speed scenarios, e.g., parking assistance.
2. *Short Range* (30 m). Forward-looking cameras are used for lane departure warning, traffic sign recognition, and backward cameras are for parking assistance. Short-range MMW radars (SRR) serve for blind spot detection and cross traffic alert.
3. *Medium Range* (80 – 160 m). LiDAR and medium-range MMW radars (MRR) operate in medium range and assist collision avoidance and pedestrian detection.
4. *Long Range* (250 m). Long-range MMW radars (LRR) are designed for Adaptive Cruise Control (ACC) at high speeds.

Because the physical principles underlying these technologies vary, their operation ranges are different as well. We emphasize the major differences of these technologies below.

Physical principle. On-board vehicle sensors for detecting barriers and road condition utilize three types of waves. Both LiDAR and cameras rely on lights (i.e., infrared and visible light) to recognize objects. In comparison, ultrasonic sensors detect obstacles by transmitting and receiving ultrasound, which is one type of mechanical waves with their frequency beyond human hearing ranges. MMW radars rely on millimeter waves, a band of electromagnetic wave whose frequency is much lower than light yet much higher than well-known radio frequency range (e.g., 2.4 GHz). Because each type of sensors rely on a distinct underlying principle, various methods and equipment have to be utilized to attack each type of sensors.

Sensor market penertration. We first try to understand the cost for each type of sensors, because the costs of manufacturing sensors determine their market shares. The costs of sensors from low to high are as follows: ultrasonic sensor, camera, radar, and LiDAR. Naturally, ultrasonic sensors have been widely deployed on modern vehicles for parking assistance systems, and the rest of sensors are reserved for high-end features. Cost-performance trade-off is perhaps the reason that car manufacturers (e.g., Tesla) abandon LiDAR [8], but self-driving prototype developers (e.g., Google [7] and Stanford [24]) tend to utilize every possible sensor to cover a large observation range.

Since not all manufacturers utilize LiDAR, we examine the rest three types of sensors that have been widely applied on existing vehicles for driver assistance systems. To the best of our knowledge, the security vulnerabilities of automotive ultrasonic sensors and MMW radars on automobiles have never been examined in practice before. Perhaps the most relevant work is done by Petit [19] on LiDAR. Both Petit's work and ours are complementary and two works in combination provide a comprehensive view on the security issues of sensors in self-driving vehicles. Apart from in-lab studies on stand-alone sensors, we carry out outdoor experiments on vehicles in this work. Note that Tesla model S cars employ all three sensors in their Autopilot systems and thus most of our work involves testing on a Tesla model S vehicle.

4. ATTACK OVERVIEW

In this section, we provide an overview of the attacks by specifying the attack assumption and introducing the basic ideas of attacks. Then, we discuss our experiment procedures.

4.1 Threat Model

Sensor Assessment. We assume that attackers have prior knowledge of the underlying principles of sensors, and have budges and access to obtain targeted sensors for further study so that they can acquire the parameters of sensor designs (e.g., operational frequency, bandwidth, duty cycles, packet format) or explore vulnerabilities of sensors. Attackers are proficient with hardware design, and able to exploit off-the-shelf hardware to accomplish their attacks.

Contactless. Except from doing prior sensor studies, attacks cannot directly contact the victim automobiles nor to the on-board sensors. Thus, the attackers have to remain outside of the vehicle to carry out the attacks, and no physical alteration or damage can be made to the targets.

Attack Scope. Although various attack surfaces have been disclosed (e.g., via cellular networks), we focus on attacking the automated driving systems and ADAS by falsifying the sensors' output. We note that the final behavior of the automobiles depends both how sensors cope with attacks and how the entire ADAS or self-driving systems handle anomaly. Thus, attackers exploit both sensors and self-driving systems to induce malfunction of automobiles, and the same attacks may result in distinct behaviors for different vehicle models.

4.2 Attack Model

We study three types of sensors: ultrasonic sensors, MMW radars, and cameras. Each operates on a distinct physical principle and requires different equipment for attacks. Nevertheless, we summarize the shared features and basic ideas of our attacks for all three types below.

4.2.1 Sensor Attacks

In this work, we refer to the attacks against sensors as *sensor attacks*. Unlike the traditional cyber attacks that alter digital information or invade systems exploiting digital channels (e.g., network interfaces, file systems, memories), sensor attacks take advantages of the analog sensing channels, i.e., utilizing the underlying physics principles of sensing to disrupt or manipulate the analog sensor measurements. Since sensor inputs play a fundamental role in a control system and are not expected to be intentionally modified, the end systems may not be able to cope with malicious sensor inputs and can result in unexpected consequences. For instance, loud acoustic signals can cause a gyroscope on a drone to report falsified data and result in a crash [22].

To understand the feat of sensors attacks against all types of sensors, we ask two questions.

1. What is required for sensor attacks? For sensors relying on distinctive physical principles, different equipment is required for attacks. For instance, we exploit ultrasound transceivers against ultrasonic sensors, radio frequency (RF) transceivers against MMW radars, and lasers against cameras. Regardless of whether ultrasound transceivers, RF transmitters, or lasers are used, they require no physical contact with the targeted sensors, and therefore render the attacks contactless.

2. What is the effective attack range? Sensor attacks are contactless, and the longer the effective range of the attacks are, the more practical the attackers are. Note that the effective range of the attacks relies on the operational range of sensors and the transmission power of various transceivers, which is limited by the budge and device constraints. The goal of the work is to validate the feasibility of the attacks, and we did not focus on intentionally maximizing the transmission power and the reported effective attack range serves as a reference. In practice, a motivated attacker can increase the transmission power and boot the effective attack range.

4.2.2 Basic Idea

All three types of sensors measure the echos reflected by obstacles, i.e., ultrasounds, MMWs, or visible lights, and two types of attacks are feasible: jamming, whereby attackers inject noises to interfere with sensors, and spoofing, whereby carefully crafted signals are injected so that they appear to bounce from non-existing obstacles and hide real ones.

1. Jamming Attacks. Injecting the same type but a stronger signal will interfere with real ones that reflect from obstacles and may cause malfunction. Sensors are typically designed to tolerant benign ambient noises, and do not expect strong interference. It is unclear whether sensors can detect objects in the presence of jamming. In cases that interference is so strong that it causes DoS attacks, it is unclear whether the sensors and the automobiles will fail gracefully and do not cause fatal accidents.

2. Spoofing Attacks. Unlike jamming attacks where arbitrary signals suffice the attacks, spoofing attacks involve emitting carefully crafted signals (e.g., ultrasonic pulses, radio chirps) that are similar to the real signals transmitted by sensors, i.e., with the same frequency, modulation, etc. As a result, the sensors interpret the spoofed signals the same way as the real ones and are deceived to detect a non-existing obstacle. Carefully adjusting the timing of the spoofed signals can 'create' a fake obstacle in various locations. Details will be discussed in Section 5, 6, and 7. We note that the spoofing attacks are feasible for ultrasonic sensors and MMW radars, but are extremely difficult for cameras via sensor attacks.

4.2.3 *Experiment Steps*

To understand the security of automotive sensors and the end systems, we aim at finding answers to the following questions.

1. How are sensors in existing automobiles designed? To obtain the details of signals, we examine stand-alone sensors in laboratory setup.

2. Are these sensors vulnerable to jamming attacks and spoofing attacks?

3. Do automobiles treat abnormal sensory data under jamming attacks or spoofing attacks properly? To understand how resilient automotive sensors and automobiles are against both attacks, we study both attacks on vehicles, including those with Autopilot functions.

4. If sensors do not operate gracefully under attacks, what defense mechanisms can be adopted to cope with sensor attacks?

In the following sections, we will answer the aforementioned questions, and illustrate experimental attacks on automotive ultrasonic sensors, MMW radars, and cameras in details.

5. ATTACKING ULTRASONIC SENSORS

Ultrasonic-based parking assistance systems were first introduced in the European market in the early 1990s. Such a system monitors the front and rear of a vehicle, and warn the driver if obstacles in the vicinity of the vehicle can cause collisions. Recently, this technology has been implemented to assist advanced functionalities like semiautomatic parking assistance, fully automatic parking, parking space detection, and Tesla's new summon feature (parking with driver outside the vehicle) [26], since ultrasonic sensors can help to probe invisible parking spaces and to park vehicles easily, quickly, and safely [11].

Besides automotive application, ultrasonic sensors have been used in many other fields, such as submarines, medical diagnostics, testing materials, and distance measurement in manufacturing and robot technologies [2, 13, 28]. We believe that the insight gained in this study can influence more than just the automotive industry.

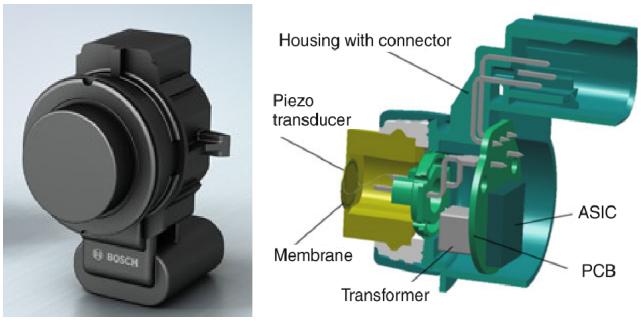

Figure 2: Appearance and cross-section of an ultrasonic sensor from Bosch.

In this section, we will first introduce the fundamentals of ultrasonic sensors, and then we present the attack methods. Finally, we show results acquired in the laboratories and outdoors. By building a DIY ultrasonic jammer using an Arduino board, we managed to launch jamming and spoofing attacks on ultrasonic sensors, and tested on several popular car models, including a Tesla Model S. Our experiments demonstrated the following attacks.

- Jamming attacks can prevent ultrasonic sensors from detecting objects, and cause collisions. In self-parking and summon mode, the Tesla model S car will ignore obstacles and crash them when jamming attack is in action.

- Spoofing attacks can manipulate the sensor measurements, and make autos to display a pseudo-obstacle.

- Acoustic cancellation is possible in theory, though sophisticated hardware and algorithms are required.

5.1 Ultrasonic Sensors

Distance measurement using ultrasound is widely used because the relatively low propagation speeds of sounds make the hardware low cost and measurement more accurate than using radio waves. To measure the distance to an object, an ultrasonic sensor emits ultrasonic pulses, and measure the time that it takes to receive echoes reflected from obstacles. The distance to the nearest obstacle is calculated based on the propagation time (time-of-flight, TOF) of the first echo pulse according to the equation

$$d = 0.5 \cdot t_e \cdot c \tag{1}$$

with t_e being propagation time of ultrasonic echoes, and c the velocity of sound in air (i.e., approximately 340 m/s). Furthermore, utilizing *trilateration* of multiple distance measurements from neighboring sensors, the objects can be localized.

Components. The sensor consists of a plastic housing with integrated plug-in connection, an ultrasonic transducer, and a printed circuit board with the electronic circuitry to transmit, receive, and process the signals, see Figure 2.

Piezoelectric Effect. A transducer inside an ultrasonic sensor emits and receives ultrasounds, in the same way as the ones for creating and receiving audible sounds (a.k.a., microphones and speakers). In the automobile industry, most

ultrasonic sensors utilize piezoelectric crystals [17], which can convert electric charges into mechanical vibrations and vice versa. If a voltage is applied at the electrodes on two sides of a piezoelectric crystal, a mechanical deformation results and generates acoustic waves. Vice versa, an incoming acoustic wave creates oscillations of the crystal, which generate an alternating voltage at the electrodes that will be amplified and digitized.

Distance Measurement. When a sensor receives a command from the ECU to transmit, its circuit excites the membrane with square waves (for approx. 300 μs) at its resonance frequency (40 – 50 kHz), resulting in sensor's vibrating and emitting ultrasound pulses. Note that a transducer cannot listen while transmitting. Even after it stops transmitting, it cannot detect echoes immediately because it takes time to stop oscillation (approx. 700 μs), and such a time is also known as the ring-down time. Because of the ring-down time, ultrasonic sensors cannot detect objects in their close vicinity and has what known as the *blanking distance*. Once rested, the membrane can be stimulated to vibrate again by the echoes reflected from obstacles. These vibrations are converted by the piezoelectric crystal to an analog signal, which is then amplified, filtered, digitized, and compared to a threshold to determine the reception of echoes. Finally, the sensor transmits the time-of-flight to the ECU for further distance calculation.

Frequency. In automotive parking aid systems, ultrasonic transducers typically operate in a frequency band between 40 and 50 kHz. This has been proved as the best trade-off between acoustical performance (sensitivity and range) and robustness against ambient noises of the transducer with the following reason. Compared with 40 – 50 kHz Higher frequencies lead to lower echo amplitudes because of higher attenuation of the airborne sounds, whereas for lower frequencies the proportion of interfering sound is larger [18].

To understand the impact of jamming on on-board ultrasonic sensors, we built an attack system that can generate ultrasounds in the same frequency band as the ones of automotive sensors and tested on cars.

5.2 Jamming Attack

Jamming attacks generate ultrasonic noises that induce continuous vibration on the sensor membrane, and render distance measurement impossible. The goal is to induce failure of obstacle detection, which may cause collisions during parking or false maneuvering.

5.2.1 Jamming

A jamming attack continuously emits ultrasounds towards a sensor to lower the SNR of the echo signal, as shown in Figure 5. To realize a jamming attack, we consider the following factors.

Resonant Frequency. From marketing materials, we learn that ultrasonic sensors for parking assistance generally operate on the frequencies between 40 kHz and 50 kHz. From our measurement on several car models, it turns out that the operation frequency appears to be near 50 kHz. Ultrasonic transducers operate around a narrow band centered at their resonant frequencies, which are determined by the diameters of the piezoceramics. Since ultrasonic sensors exhibit high sensitivity within several kHz of the resonant frequencies, the jamming transducers that operate in the same frequency band can result in an effective jamming disturbance, which in our case is 50 kHz. Unfortunately, 50-kHz transducers are unavailable on the market, and we have to use the off-the-shelf 40 kHz transducers, which turned out to be effective in jamming and we believe that the effective distance will be expended if matching transducers are available.

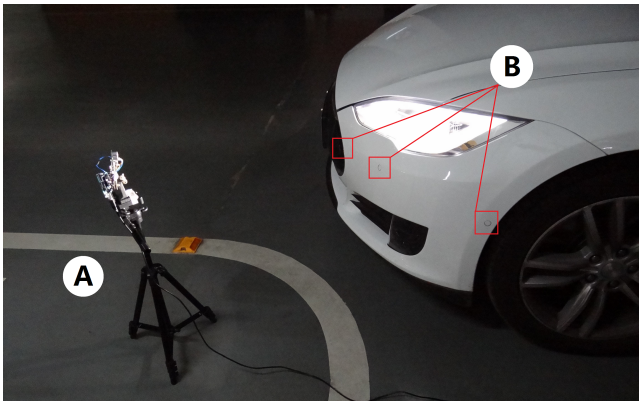

Figure 3: Setup of ultrasound experiment on Tesla Model S. A is the jammer, B is 3 ultrasonic sensors on the left-front bumper.

Emitting Ultrasound. Applying alternating voltage on piezoelectric crystals generates acoustic waves, and the frequency of the AC input signals determines the oscillation frequency and therefore the frequency of acoustic waves. By applying a 40 kHz square wave to the transducer, we are able to generate ultrasounds of 40 kHz. The same principle works for other frequencies as long as the hardware (e.g., speakers) has the desired resonant frequency.

Equipment. To generate controllable square waves at 40 kHz, we utilized off-the-shelf hardware, e.g., Arduino board [1] and function generator. Arduino can output a square wave of selected frequencies on the digital I/O pins using a built-in function called `Tone()`, which is mainly used to generate tones for speakers. Due to the low-cost nature, Arduino cannot drive a perfect square wave without any frequency jitters. Nevertheless, the generated square waves are sufficient for driving jamming ultrasounds. A function generator can output signals with better frequency performance and higher amplitude, and achieves longer jamming distance.

Voltage Level. The amplitudes of sounds created by piezoelectric crystals rely on the voltage level, and vice versa. Since the effectiveness of jamming ultrasounds is determined by their transmission power, the effective attack distance is decided by the applied voltages. In our experiments, we use two type of equipment. Arduino can output square waves with 5 volts maximum, and a function generator can generate a voltage up to 20 volts. The ultrasonic sensors that we obtained can take up to 70 volts, and we believe that the effective attack range can go beyond what we observed as long as a 70-volt square wave generator can be acquired.

5.2.2 Results

We have validated jamming attacks on the following three types of scenarios: (1) stand-alone ultrasonic sensors, (2) cars with parking assistance, and (3) a Tesla Model S with self parking and summon. In all experiments, a real obstacle

existed and it can be detected by the sensor when no attack is in progress.

Ultrasonic Sensors. First, we tested 8 models of standalone ultrasonic sensors in the laboratory. Six of them are individual ultrasonic ranging modules, one of them is an aftermarket vehicular sensor, and the other is an OEM parking assistance system consisting of one ECU and four sensors. Under jamming attacks, we observed two types of sensor outputs: one is *Zero distance*, while the other is *Maximum distance*. *Zero distance* means that the sensors detect an obstacle within 10 cm, and *Maximum distance* indicates that nothing is detected. They are the result of two types of sensor designs to process the measured echoes. For *Zero distance*, a sensor will consider the existence of an obstacle if the amplitude of received ultrasounds is larger than a pre-defined threshold. As soon as the sensor passes the ring-down period, it will receive the loud jamming signal and consider it as the echoes from obstacles, resulting in zero distance. For *Maximum distance*, the sensors is designed to suppress ambient noises by adjusting its threshold accordingly. A high level of ambient noises maps to a high threshold. Our jamming signal is recognized as noises because it exists throughout the entire cycle. To suppress the ambient noise, the sensor raises the threshold so that the amplitude of the legitimate echoes is smaller than the threshold and hence maximum distance.

Cars with Parking Assistance. Next, we examined a few vehicles with driver assistance systems, which include popular models from Audi, Volkswagen, Tesla, and Ford. The driver assistance systems on these cars differ in terms of sensor brands and ECUs. Nevertheless, they all inform the driver obstacles by either vocal or visual display. As shown in Figure 3, the ultrasonic jammer is placed in front of the car bumpers and can be correctly detected before jamming attacks start. Once a jamming attack is launched, the vehicle can no longer detect the obstacle and no alarm is triggered (Figure 4(c)). We believe that this maps to the *Maximum distance* case and the design of these sensors aims at noise reduction. We tested the vehicle in both parking and reverse gear, and the results remain the same. Using a function generator we can launch the attacks from 10 meters away for Tesla Model S. The consequence of jamming the sensors in parking assistance systems is collision, which could be serious when pedestrians are the obstacles.

Tesla Model S with Automatic Parking. We further tested jamming attack on the self parking and summon functionalities of Tesla Model S. We were wondering whether jamming attacks can prevent automatic parking systems from detecting the obstacles reliably. Without human supervision, the aftermath of jamming attacks could be even worse than driver assistance systems. The results we observed turned out to be prominent and worrisome. When the Tesla is in self-parking or summon mode, and jamming attacks are in action, the car moving by itself will ignore obstacles and collide with them. When the jammer is drived by an Arduino, the distance between the interferer and car is 20 cm. However, the range can be increased to 1 meter with a function generator. We believe the difference is caused by the power of the jammer. When the jammer is powerful and close-by, the jamming signal is strong enough to suppress the legitimate echo signals, while when the jamming signal is not powerful enough, the attack is not successful.

(a) Normal. (b) Spoofed. (c) Jammed.

Figure 4: Tesla parking distance display at normal, being spoofed, and being jammed[1].

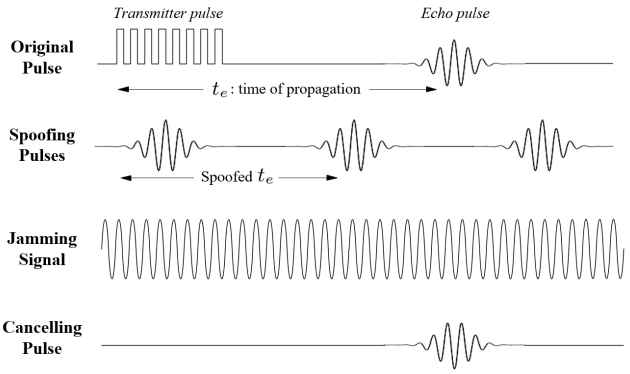

Figure 5: Illustration of all ultrasonic attacks. From up to down are original signal, spoofing signal, jamming signal, and acoustic cancellation signal. The last 3 attack signals overlay with the original signal at the sensor side.

5.3 Spoofing Attack

We use the same equipment to launch spoofing attacks, yet instead of sending ultrasonic signals all the time, we have to send ultrasonic signals at the right timing to deceive sensors, i.e., manipulation of the sensor readings, which can lead drivers to reduce their confidence on trusting sensors.

5.3.1 Spoofing

Spoofing attacks are based on the idea that when carefully crafted ultrasound pulses from adversaries can be falsely recognized as echoes from obstacles, and arrive at the sensor ahead of the real ones, then the sensor will conclude the detection of an obstacle closer than the real one. By adjusting the arriving time of spoofing pulses, an attacker can manipulate sensor readings, i.e., distance measurement. An illustration is shown in Figure 5. We note that it is possible to combine spoofing and jamming signal so that the distance can be both decreased and increased.

Setup. The experiment setup is similar to the ones for jamming attacks, except that the transducer is excited with 50 kHz square wave, which produces the same signals as the real ones.

Excitation Time. Although a legitimate probing pulse lasts for 300 μs on vehicles, an excitation time of 200 – 500 μs generally suffices for spoofing the sensors.

[1]This is a default display of tire pressure. It pops out every time we do ultrasonic jamming, and disappears when we stop. Anyway, NO distance information can be displayed during jamming.

Timing. Transmitting at the right timing is non-trivial for a successful spoofing attack. Unlike LiDAR, ultrasonic sensors only care about the nearest obstacle. This means only the first justifiable echo will be processed, other following echoes will be ignored. Thus the spoofing pulse has to arrive ahead of the real ones to be effective. Here we define the *Effective Slot* for spoofing attacks, as the time slot between the end of the sensor probing pulse and the start of the first real echo received. The spoofing pulse must reside within the *Effective Slot*, the length of which depends on the obstacle distance. In addition, automotive ultrasonic sensors are expected to detect an obstacle at most 2 m away, therefore the spoofing pulse has to be received within 11.7 ms (a maximum for the *Effective Slot*) from when the probing signals are transmitted. Given that a car may send out probing pulses every 100 ms (or more), an attacker has less than 11.7% window to accomplish the attack. Moreover, in reality a car will not transmit probing pulses at a fixed period, due to intentional jittering or asynchronous cycles [18]. Hence injecting spoofing pulses blindly or based on prediction will probably fail. To overcome this issue, we inject the spoofing pulses every several milliseconds. It may cause unstable spoofed sensor readings, but can greatly increase the probability of successful injecting in the *Effective Slot*.

5.3.2 Results

As mentioned above, a successful spoofing attack depends on the timing of injection, as well as the length and cycle of spoofing pulse. By trial and error we are able to find a set of parameters that can produce interesting sensor outputs, such as abrupt change, steady oscillation between near and far, and altering around a certain reading, as shown in Figure 4(b). In most cases, the sensor readings are just disturbed randomly. When there is no obstacle in the detection range at all, spoofing attack can cause the display of pseudo-obstacles.

5.4 Acoustic Quieting

Instead of jamming attack, an alternative way to hide obstacles from ultrasonic sensors is to eliminate echoes. This approach of *Acoustic Quieting* has been well researched [4, 5, 14], and well developed for miliary submarines to stay stealth [10, 29]. Methods include silent running, hull coatings that reduce active sonar responses, and hydrodynamic hull design that reduces noises and active sonar responses. We propose two similar methods of acoustic quieting for vehicles.

Cloaking. Sound absorbing materials (e.g., plastic foams) are hardly seen by the ultrasonic parking system. For people wearing sound absorbing cloths (e.g., woman with a furcoat), the system has a shorter detection range. Our idea is to cover the obstacle or human with deadening like sound absorbing foam. From our experiments, damping foams can eliminate the majority of the returning echoes, and easily hide obstacle or human from sensors.

Acoustic Cancellation. *Active Noise Control* (ANC), also known as noise cancellation, or Active Noise Reduction (ANR), is a method for reducing unwanted sound by the addition of a second sound specifically designed to cancel the first [6]. Helicopter pilots rely on this technology to speak on the radio; it is also implemented on many high-end headphones. Though originally designed for cancelling low frequency noises, we believe this method can also be applied to cancel ultrasound pulses from vehicular sensors, because the frequency is fixed and the time pattern is predictable. Note that the cancelling pulse in Figure 5 is in reverse phase to the original one. We have done preliminary experiments that proved the feasibility of canceling ultrasound by minor phase and amplitude adjustment. We refer intentional readers to search for high-speed hardware for vehicular ultrasound cancellation.

6. ATTACKING MMW RADARS

Radar (Radio Detection and Ranging) originates from the military technology since the Second World War, and has been bound to military applications for a long time. The first vehicle with Radar for adaptive cruise control was made available until 1998. Five years later, this technology was boosted due to the development of automatic emergency brakes and lane changing assistance. Automotive radars have different requirements and solutions compared to military applications, such as smaller distance, lower Doppler frequency, higher multitarget capability, smaller sizes, and significantly lower cost [9, 21]. Although various ranges available, a Medium Range Radar (MRR) is mounted in the front grill on Tesla Model S to support many of the Autopilot functions, e.g., front collision avoidance and traffic-aware cruise control.

In this section, we will present our study on the Radar and Autopilot system in Tesla Model S. From a signal analyzer we were able to identify the frequency band, modulation scheme, and waveform pattern of the Tesla Radar. Then we tried to jam and spoof the radar system with electromagnetic waves in the same frequency band generated by a signal generator. Our experiments show that automotive MMW Radars can suffer from electromagnetic jamming and spoofing. We will demonstrate the following:

- Jamming attacks can make detected objects disappear from the Radar and Autopilot system.
- Spoofing attacks can alter the measured distance of the obstacle.

6.1 MMW Radars

This section presents an overview on the basic principles of Radar telecommunication technology in layman's terms.

Basic Principle. Radars work on the basic principle of active probing, i.e., transmitting electromagnetic waves for probing, receiving the electromagnetic echo that is bounced back from an object, and measuring the echo's parameters, including but not limited to the time-of-flight. However, due to the way faster propagation speed of electromagnetic waves ($3 \times 10^8 m/s$), the methods used for ultrasonic sensors are no longer technically feasible. The emitted electromagnetic waves must be given an identifier for recognition and a time reference for the measurement of time-of-flight, which is referred to as modulation. At the receiver, demodulation is performed. The waveform can be described as a harmonic wave function in a general form:

$$u_t(t) = A_t \cdot \cos(2\pi f_0 t + \varphi_0) \tag{2}$$

Modulation is therefore possible with three variables: amplitude A, frequency f, and phase φ. Amplitude modulation is basically pulse modulation; frequency modulation

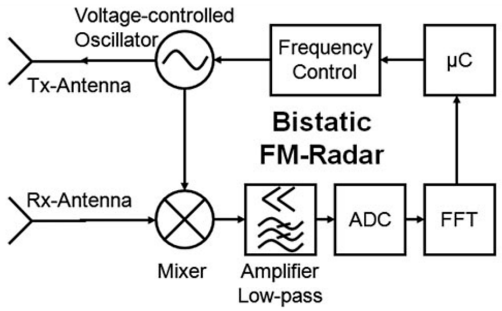

Figure 6: Block diagram of a bistatic Radar with frequency modulation [30].

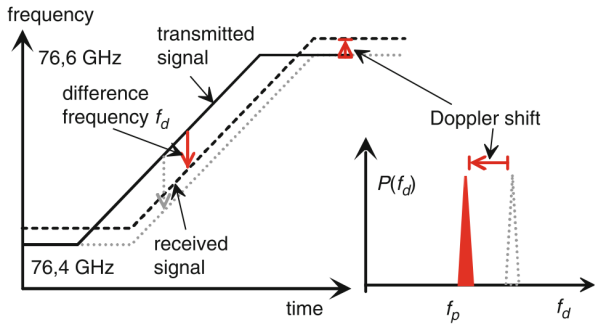

Figure 7: Spectral display of FMCW with a positive ramp for an approaching object [30].

includes *Frequency Shift Keying* (FSK), *Frequency Modulated Shift Keying* (FMSK), *Frequency Modulated Continuous Wave* (FMCW), and *Chirp Sequence Modulation*. In the scope of this paper, frequency modulation and FMCW especially are introduced, because they are widely used in reality.

Frequency Modulation. In frequency modulation, the frequency f_0 varies as a function of time. Figure 6 shows a basic structure of FM radar. The instantaneous frequency is controlled by a voltage-controlled oscillator (VCO) which enables the desired modulation via a control loop (e.g., phase-locked loop (PLL)). The received signal is then mixed[2] with the signal currently being transmitted, amplified, filtered, digitized, and converted to the frequency domain for further processing.

FMCW. Frequency modulated continuous wave is a frequently used modulation scheme for automotive radars. As shown in Figure 7, the instantaneous frequency is continuously changed in the form of a linear ramp. With known slope m_ω, the measurement of time-of-flight can be converted to the measurement of frequency difference f_d, which can be measured easily by signal mixing. The relative speed can be further calculated from the Doppler shift. By means of additional ramps with different slopes m_ω, the ambiguity of linear combination can be resolved for a small number of objects.

Doppler Effect. If an object moves relative to the radar, the reflected electromagnetic wave will undergo a frequency shift, which is described as the Doppler Effect. Accordingly, the frequency shift can be used to measure the relative velocity.

Frequency Bands. There are currently four bands available in road traffic (24.0 − 24.25 GHz, 76 − 77 GHz, and 77 − 81 GHz in addition to a UWB band of 21.65 − 26.65 GHz suitable for close range). The 76.5 GHz range, which is exclusive for automotive Radar and available worldwide, dominates at present. The 24 GHZ range has also claimed a large share of the market, especially for medium-range and close-range applications.

Attenuation. Atmospheric attenuation is below 1 dB/km at 76.5 GHz, and therefore only 0.3 dB for the return path to a target 150 m away. However, heavy rain with big raindrops that achieve the magnitude of the wave length (3.9 mm) will result in serious attenuation, and leads to significant range reduction. In addition, heavy rain results in an increased interference level (clutter) and decreases the SNR, which will in turn reduce the detection range.

6.2 Signal Analysis

Knowledge of MMW frequency range, modulation, ramp pattern, etc. of the radar is required for crafting attacks, but the Radar technology used on Tesla Model S is not publicly known. Instead of tearing down the front bumper and messing with the Radar, we observed the radar spectrum and waveform, and reverse engineered the radar waveform directly, which is challenging because of the high frequency.

6.2.1 Signal Analysis

We learn that a Bosch 76 – 77 GHz MRR Radar is installed on Tesla Model S. We utilize professional equipment to analyze such a high frequency. However, normal spectrum analyzers and signal generators merely work at frequencies as high as several giga Hertz, the ones we have access to can only reach 40 – 50 GHz, therefore we use frequency multipliers and mixers to handle signals at 77 GHz.

Equipment. The following equipment have been employed for signal analysis: Keysight N9040B UXA Signal Analyzer (3 Hz – 50 GHz), DSOS804A High-Definition Oscilloscope, 89601B VSA Software, and VDI 100 GHz harmonic mixer. The mixer acts as the RF frontend and down-converts the 77 GHz signal to a lower frequency that the signal analyzer can process. An oscilloscope is connected to the signal analyzer for better observation in the time domain. VSA software is used for further signal analysis.

Experiment Setup. Figure 8 shows the setup of Radar experiments. A frequency mixer with a horn antenna is connected to the signal analyzer, from which we observe and analyze the signal. To receive Radar signal of higher amplitude, we placed the antenna 0.5 m ahead of the car and on the same horizonal level in line with the automotive Radar.[3] After switching to the Drive gear, the Radar on Tesla is powered on, which can be confirmed from the display of a virtual car in the middle of the dashboard (Figure 9(a)). This virtual car maps to the cart containing all equipment in reality. To simplify the analysis, we kept our cart and equipment still throughout the experiments.

6.2.2 Results

[2]The process of signal multiplication is described as mixing in high-frequency technology. By mixing it is possible to measure the signal at much lower frequencies.

[3]A caution of safety in doing the alignment is NOT to look at the functioning Radar closely and directly in the eyes. It will damage your eyes.

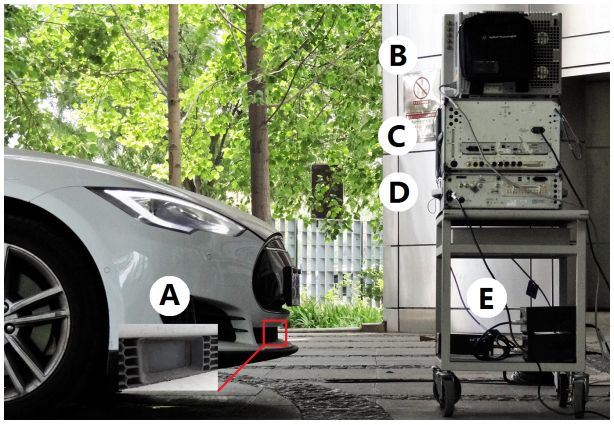

Figure 8: Setup of Radar experiments on Tesla Model S. A is automotive Radar, B is oscilloscope, C is signal analyzer, D is signal generator, E is the collection of frequency multiplier, harmonic mixer, and their power supplies. Oscilloscope, signal analyzer, and harmonic mixer are used in signal analysis. Signal generator and frequency multiplier are used in jamming and spoofing attacks.

Based on our observation and calculation, the center frequency of radar signals is around 76.65 GHz. After some correction and manual calculation, the bandwidth (ramp height) is approximately 450 MHz, which proves that the automotive radar on Tesla works within the 76 – 77 GHz band. The modulation is FMCW with 5 chirps of low ramp slope, which agrees with the technical data of Bosch MRR 4.

6.3 Jamming Attack

A straightforward idea of attacking radars is jamming them within the same frequency range, i.e., 76 – 77 GHz.

6.3.1 Jamming

In a normal scenario, to measure the distance to an obstacle, the received radar signals that are bounced back have to be sufficiently stronger than the noises. Depending on any other signal evaluation for flare suppression, the threshold typically is above the electrical noise by a factor SNR threshold of approximately 6 – 10 dB [30]. Jamming signals increases the noise level and will reduce SNR, and therefore lead to radar system failure.

Jamming Waveform. Because the Radar signal sweeps 450 MHz, jamming within this range will likely succeed. We came up with two approaches, one is jamming at a fixed frequency of 76.65 GHz, and the other is sweeping frequency within the 450 MHz range.

Equipment. Keysight N5193A UXG Agile Signal Generator (10 MHz – 40 GHz) and VDI WR10 frequency multiplier (75 – 110 GHz) are used together as an interferer that emits electromagnetic waves at 77 GHz. The signal generator generates jamming signal at 12.775 GHz, and transmit it to the frequency multiplier, which multiplies the frequency by 6 to 76.65 GHz, and emits through a horn antenna.

Experiment Setup. The setup is similar to Figure 8, except that firstly we have to ensure that an object has already been detected by the Radar system before jamming,

(a) Drive gear. (b) Autopilot. (c) Jammed.

Figure 9: Tesla dashboard display at drive gear, Autopilot, and Autopilot under radar jamming.

so that the outcome can be observed. We keep the equipment cart still and the interferer off, and drive the Tesla toward them until a virtual car is displayed on the dashboard. The final distance between the car and equipment is basically between 2 – 3 m. When the Autopilot mode is further turned on[4], the virtual car turns from black to blue (Figure 9(b)). Interestingly, in Autopilot mode the cart is easier to be recognised as a car than in the hand-driving mode. After settling the car, we turn the interferer on and off, and observe the information displayed in Tesla.

6.3.2 Results

The outcome of jamming attack is prominent. After settling the Tesla and before turning on the interferer, a car (the cart actually) is detected by the Radar system and displayed as a virtual car on the dashboard. When the interferer is turned on, the virtual car disappears from the dashboard immediately. When the interferer is turned off, the virtual car reappears, as shown in Figure 9(c). We have repeated the experiments many times, including trying different jamming waveforms and with Autopilot on and off, the same results occur every time. Jamming attack can make detected objects disappear from the Radar and the Autopilot system. Compared with the hand-driving mode, we have discovered that the effective attack range and angle in Autopilot mode is larger. We suspect that this is caused by threshold adjustment in Autopilot for tracking objects.

6.4 Spoofing Attack

By modulating signals the same way as the automotive radar, we managed to launch spoofing attacks. However, due to the low ratio of working time over idle time, signal injection at the precise time slot is unlikely to be successful as we expected. Nevertheless, by tuning the jamming ramp slope back and forth, we happened to observe periodic distance change of the virtual car displayed in Tesla.

7. ATTACKING CAMERAS

In addition to radars, LiDAR, ultrasonic sensors, GPS, and many other sensors, images from cameras are useful to acquire road signs and lanes, especially on highways and city streets where many rules and regulations are applied. Onboard camera systems handle visual recognition of the surroundings in automated driving technology. Recognition includes lane lines, traffic signs and lights, vehicles, and pedestrians. After fusing data with other sensors, the driving behavior and routes can be improved. On Tesla for example,

[4]It is possible to turn on Autopilot when Tesla is not moving, which relieves the trouble of moving experiments.

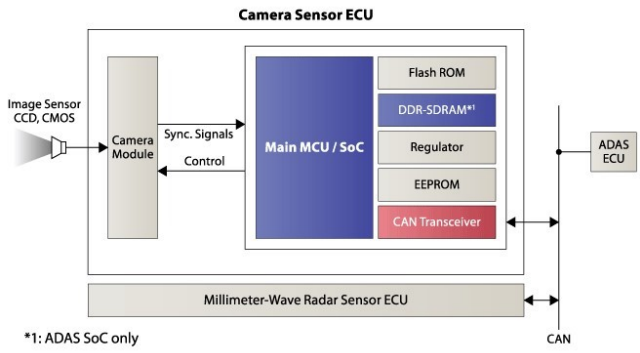

Figure 10: Forward-looking camera system block diagram [20].

a forward facing camera is used to recognize lanes and road signs. Features based on this technology include automatic lane centering and changing, lane departure warning, and speed limit display.

Cameras are passive light sensors. From our daily experience, they can be blinded or fooled in many ways. To validate the attack on vehicle cameras, we carried out blinding attacks in different scenarios, observed and recorded the camera output. This section will present the experiments on blinding the vehicle camera with lights of different wavelengths generated by off-the-shelf, low-cost light sources. Our major finding is:

- Automotive cameras do not provide enough noise reduction or protection, and thus can be blinded or permanently damaged by a strong light source.

7.1 Cameras

As shown in Figure 10, cameras collect optical data by CCD/CMOS devices through filters, generate images in the camera module, and send them to the MCU for further processing and calculation. The recognition results will be sent to the ADAS ECU from the CAN bus. ADAS processor makes driving decisions and send commands to actuators, e.g., hydraulic steering wheel and control panel. Some systems further provide the driver with video outputs on the screen for reference.

7.2 Blinding Attack

Our attack is based on the assumption that CMOS/CCD sensors can be disturbed by malicious optical inputs, and will produce unrecognizable images. The broken images will further influence the decision of ADAS unit and indirectly affect vehicle control. As a consequence, it will lead to the car's deviation, or an emergency brake, which could all possibly cause crashes.

7.2.1 *Blinding*

A common method to attack video equipment is laser blinding. Photoelectric sensors are sensitive to the intensity of light. With a peak adsorption coefficient at generally 10^3 to 10^5, most of the laser energy at the sensor can be absorbed. The time necessary for damaging photoelectric sensor is one to several orders of magnitude less than the time for harming human eyes. Under laser exposure, the surface temperature will rise rapidly due to the thermal stress caused by non-uniform temperature field. Avalanche breakdown of semiconductor materials can cause irreversible damage to the photoelectric devices. Camera exposure to laser radiation for vehicles running on the road can happen when LiDARs are nearby. LEDs can also be used to generate bright light against cameras. In our experiment, we used three kinds of light sources, i.e., LED, visible laser, and infrared LED.

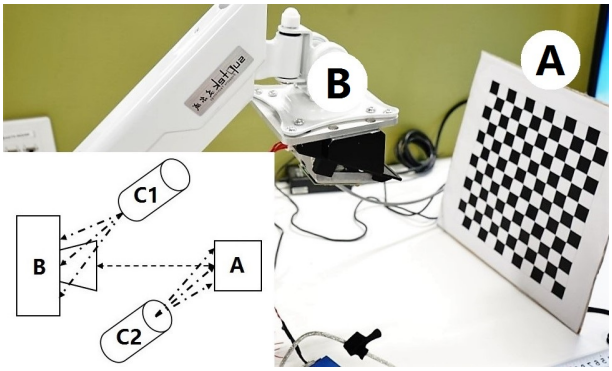

Figure 11: Setup of camera blinding experiment. A is a calibration board, B is a camera, C1 and C2 are laser emitters.

Experiment setup for blinding attack is illustrated in Figure 11. A calibration board A is positioned 1 meter in front of camera B; laser sources are either pointed at the camera or at the calibration board as C1 and C2. C1 is of 15° to the axis of A–B, and C2 of 45°. We have tested with 650nm red laser, 850 nm infrared LED spot, and LED spot of 800 mW power respectively, observed the camera image output, and measured the change of tonal distribution.

7.2.2 *Results*

LED. Aiming LED light at the calibration board leads to increased tonal value in the center area, thus information in this area can be fully concealed, and recognition will no longer be possible. Aiming LED light directly at the camera will induce significantly higher tonal values, and cause complete blindness all over the image. There is no way the camera system can acquire any visual information. The blinding time is relevant to camera refresh rate, as well as the distance between light source and camera. The results are shown in Figure 12.

Laser. Pointing laser beam at the calibration board have almost no effect on the camera. However, pointing directly toward the camera will lead to complete blindness for approximately 3 seconds, during which the recognition will be

(a) Toward board. (b) Toward camera. (c) Tonal distribution.

Figure 12: Blinding camera with LED spot.

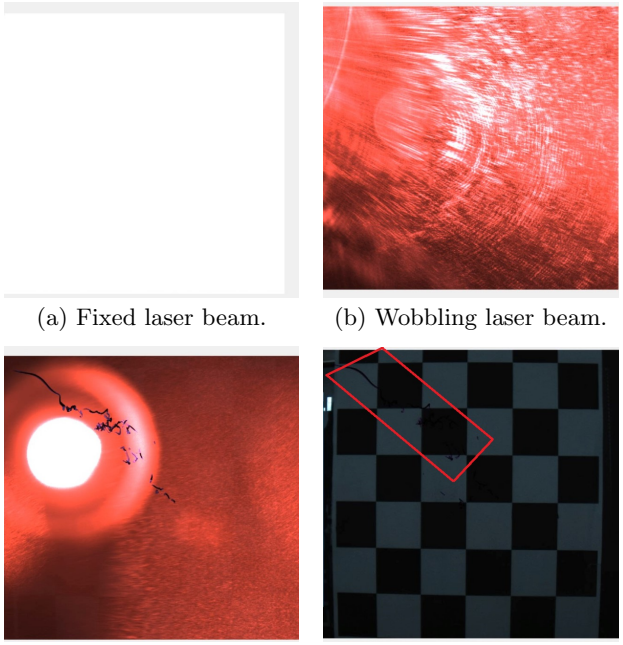

(a) Fixed laser beam. (b) Wobbling laser beam.

(c) Damage caused by laser. (d) Damage is permanent.

Figure 13: Blinding camera with confronted laser.

impossible. We further did another experiment with wobbling laser beam to emulate handhold attacks or unintentional scenarios. As shown in Figure 13(b), it can also cause failure of camera image recognition, though the tonal values are not as high due to shorter exposure time at one spot of CMOS/CCD chip.

Permanent Damage. When a laser beam is directly radiated at the camera within 0.5 meter for a few seconds, irreversible damage is caused to the CMOS/CCD chip. The black curve in Figure 13(c) is the evidence. When the laser is turned off, the curve still remains, as in Figure 13(d). Therefore, the damage is permanent and irreversible and can only be fixed by replacing the CMOS/CCD component. Unintentional damage of this kind can possibly be caused by nearby laser radars.

Infrared LED. No effect on the camera has been observed by pointing the infrared LED spot either at the camera or board. We assume it is due to narrow frequency band of filters on the camera, which is a sign of good hardware quality.

8. CONCLUSIONS

This paper exhibits that sensor resilience to attacks is an important aspect of the security and safety of autonomous vehicles. We have examined three types of sensors that Automated Driving Systems rely on and have been deployed on Tesla vehicles with Autopilot, i.e., ultrasonic sensors, Millimeter Wave Radars, and cameras. All of them are vulnerable to intentional attacks and caused malfunction in the automotive system, all of which could impair the safety of self-driving cars and require an improved design to alleviate the issues.

9. ACKNOWLEDGEMENTS

We thank Dr. Xin Bi and Keysight Open Laboratory & Solution Center in Beijing for their professional support and for providing access to the Radar equipment. We thank Xpwn Team for participating in the ultrasound research. We thank our colleagues, Weibin Jia, Zhou Zhuang, Guoming Zhang at Zhejiang University USSLab, and Bin Guo at Qihoo 360 ADLAB for their great support in the experiments.

10. REFERENCES

[1] Arduino. Arduino and Genuino Project. https://www.arduino.cc/. Accessed: 2016-07-05.

[2] L. Bergmann. The ultrasound and its application in science and technology. 1954.

[3] S. Checkoway, D. Mccoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive Experimental Analyses of Automotive Attack Surfaces. *System*, pages 1–6, 2011.

[4] H. Chen and C. Chan. Acoustic cloaking in three dimensions using acoustic metamaterials. *Applied physics letters*, 91(18):183518, 2007.

[5] S. A. Cummer and D. Schurig. One path to acoustic cloaking. *New Journal of Physics*, 9(3):45, 2007.

[6] S. J. Elliott and P. A. Nelson. Active noise control. *IEEE signal processing magazine*, 10(4):12–35, 1993.

[7] Google. Google Self-Driving Car Project. https://www.google.com/selfdrivingcar/. Accessed: 2016-07-06.

[8] S. Hall. Elon Musk says that the LIDAR Google uses in its self-driving car 'doesn't make sense in a car context'. http://9to5google.com/2015/10/16/. Accessed: 2016-07-06.

[9] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt. Millimeter-wave technology for automotive radar sensors in the 77 ghz frequency band. *IEEE Transactions on Microwave Theory and Techniques*, 60(3):845–860, 2012.

[10] L. He. Development of submarine acoustic stealth technology. *Ship Science and Technology*, 28(s2):9–17, 2006.

[11] R. Katzwinkel, R. Auer, S. Brosig, M. Rohlfs, V. Schöning, F. Schroven, F. Schwitters, and U. Wuttke. Einparkassistenz. In *Handbuch Fahrerassistenzsysteme*, pages 471–477. Springer, 2012.

[12] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental security analysis of a modern automobile. *Proceedings - IEEE Symposium on Security and Privacy*, pages 447–462, 2010.

[13] H. Kuttruff. *Ultrasonics: Fundamentals and applications*. Springer Science & Business Media, 2012.

[14] J. Li and J. Pendry. Hiding under the carpet: a new strategy for cloaking. *Physical Review Letters*, 101(20):203901, 2008.

[15] C. Miller and C. Valasek. A Survey of Remote Automotive Attack Surfaces. *Defcon 22*, 2014.

[16] C. Miller and C. Valasek. Remote Exploitation of an Unaltered Passenger Vehicle. *Blackhat USA*, 2015:1–91, 2015.

[17] M. Noll and P. Rapps. Ultraschallsensorik. In *Handbuch Fahrerassistenzsysteme*, pages 110–122. Springer, 2012.

[18] M. Noll and P. Rapps. Ultrasonic sensors for a k44das. In *Handbook of Driver Assistance Systems: Basic Information, Components and Systems for Active Safety and Comfort*, pages 303–323. Springer, 2016.

[19] J. Petit, B. Stottelaar, M. Feiri, and F. Kargl. Remote Attacks on Automated Vehicles Sensors: Experiments on Camera and LiDAR. *Blackhat.com*, pages 1–13, 2015.

[20] Renesas. Front Detection. https://www.renesas.com/zh-cn/solutions/automotive/adas/front.html. Accessed: 2016-07-07.

[21] M. Skolnik. An introduction and overview of radar. *Radar Handbook*, 3, 2008.

[22] Y. Son, H. Shin, D. Kim, Y. Park, J. Noh, K. Choi, J. Choi, and Y. Kim. Rocking drones with intentional sound noise on gyroscopic sensors. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 881–896, 2015.

[23] R. Staszewski and H. Estl. Making cars safer through technology innovation. *White Paper by Texas Instruments Incorporated*, 2013.

[24] S. A. D. Team. Welcome. http://driving.stanford.edu/. Accessed: 2016-07-06.

[25] Tesla. A tragic loss. https://www.teslamotors.com/blog/tragic-loss, June 2016.

[26] Tesla Motors. *Tesla Model S Software Release Notes v7.1*, 2016.

[27] C. Valasek and C. Miller. Adventures in Automotive Networks and Control Units. *Technical White Paper*, page 99, 2013.

[28] J. Waanders. Piezoelectric ceramics-properties and applications. philips components. *Marketing Communications*, 1991.

[29] Wikipedia. Teardrop hull. https://en.wikipedia.org/wiki/Teardrop_hull. Accessed: 2016-07-06.

[30] H. Winner. Automotive radar. In *Handbook of Driver Assistance Systems: Basic Information, Components and Systems for Active Safety and Comfort*, pages 325–403. Springer, 2016.

[31] M. Wolf, A. Weimerskirch, and C. Paar. Security in Automotive Bus Systems. *Proceedings of the Workshop on Embedded Security in Cars*, pages 1–13, 2004.

[32] M. Wolf, A. Weimerskirch, and T. Wollinger. State of the art: Embedding security in vehicles. *EURASIP Journal on Embedded Systems*, 2007(1):1–16, 2007.